

A Bayes interpretation of stacking for \mathcal{M} -complete and \mathcal{M} -open settings

Tri Le¹, Bertrand Clarke¹

Abstract

In \mathcal{M} -open problems where no true model can be conceptualized, it is common to back off from modeling and merely seek good prediction. Even in \mathcal{M} -complete problems, taking a predictive approach can be very useful. Stacking is a model averaging procedure that gives a composite predictor by combining individual predictors from a list of models using weights that optimize a cross-validation criterion. We show that the stacking weights also asymptotically minimize a posterior expected loss. Hence we formally provide a Bayesian justification for cross-validation. Often the weights are constrained to be positive and sum to one. For greater generality, we omit the positivity constraint and relax the ‘sum to one’ constraint.

A key question is ‘What predictors should be in the average?’ We first verify that the stacking error depends only on the span of the models. Then we propose using bootstrap samples from the data to generate empirical basis elements that can be used to form models. We use this in two computed examples to give stacking predictors that are (i) data driven, (ii) optimal with respect to the number of component predictors, and (iii) optimal with respect to the weight each predictor gets.

Keywords: stacking, cross-validation, Bayes action, prediction, problem classes, optimization constraints

Email addresses: `t1e20@unl.edu` (Tri Le), `bclarke3@unl.edu` (Bertrand Clarke)

¹Department of Statistics, University of Nebraska-Lincoln.

1. Introduction

Stacking is a model averaging procedure for generating predictions first introduced by Wolpert (1992). The basic idea is that if J candidate signal plus noise models of the form $Y = f_j(x) + \epsilon$ for $j = 1, \dots, J$ are available then they can be usefully combined to give the predictor

$$\hat{Y}(x) = \sum_{j=1}^J \hat{w}_j \hat{f}_j(x),$$

where \hat{f}_j is an estimate of f_j . Usually, $f_j(x) = f_j(x, \beta_j)$ so $\hat{f}_j(x) = f_j(x, \hat{\beta}_j)$ where $\hat{\beta}_j$ is an estimate of β_j . The weights $\hat{w} = (\hat{w}_1, \dots, \hat{w}_J)$ satisfy

$$\hat{w} = \arg \min_w \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{f}_{j,-i}(x_i) \right)^2 \quad (1)$$

where $\hat{f}_{j,-i}$ is the estimate of f_j using the $n - 1$ of the n data points by dropping the i -th one, i.e., $(x_1, y_1), \dots, (x_{i-1}, y_{i-1}), (x_{i+1}, y_{i+1}), \dots, (x_n, y_n)$. The Y_i 's are assumed independent and the x_i 's are deterministic design points. Often the w_j 's are assumed to be non-negative and sum to one. The properties of stacking as a predictor have been explored in numerous contexts such as regression Breiman (1996), Clarke (2003), Sill et al. (2009), classification and distance learning Ting and Witten (1999), Ozay and Vural (2012), density estimation Smyth and Wolpert (1999), and estimating bagging's error rate Rokach (2010), Wolpert and Macready (1999).

These earlier contributions treated stacking as a frequentist procedure. However, more recently, Clyde and Iversen (2013) brought stacking into the Bayesian paradigm. They recalled the tripartite partition of statistical problems into three classes namely \mathcal{M} -closed, \mathcal{M} -complete, \mathcal{M} -open, see Bernardo and Smith (2000), and suggested that outside the \mathcal{M} -closed setting the posterior risk could be approximated by a cross-validation error (for the same loss function). Hence the action minimizing the posterior risk could be approximated by the stacking predictor that minimizes (1). More precisely, given models M_j for $j = 1, \dots, J$, a loss function ℓ , a vector of responses $\mathbf{Y} = \mathbf{y} = (y_1, \dots, y_n)$, and an element

$a(\mathbf{y})$ in the action appropriate for a collection of models, say \mathcal{M} , Clyde and Iversen (2013) used

$$\int \ell(y_{n+1}, a(\mathbf{y})) p(y_{n+1} | \mathbf{y}) dy_{n+1} \approx \frac{1}{n} \sum_{i=1}^n \ell(y_i, a(\mathbf{y}_{-i})) \quad (2)$$

in an \mathcal{M} -open context, where y_{n+1} represents a future outcome at a future design point x_{n+1} , \mathbf{y}_{-i} is the data vector \mathbf{y} with the i -th entry deleted, and $p(\cdot | \mathbf{y})$ is the predictive distribution for Y_{n+1} . Here and elsewhere, the design points x_1, \dots, x_{n+1} are suppressed in the notation unless consideration of them is essential for a step in a proof. Hence, Clyde and Iversen (2013) observed that minimizing the left hand side of (2) over $a(\mathbf{y})$ and the right hand side over $a(\mathbf{y}_{-i})$ leads to two actions that are asymptotically identical. Otherwise put, the stacking predictor is the asymptotic Bayes action for \mathcal{M} -complete problems. It is not the Bayes action in the \mathcal{M} -open case because the mode of convergence is undefined. Nevertheless, Clyde and Iversen (2013) used (2) in an \mathcal{M} -open context to good effect. It should be noted that (2) seems to have been initially conjectured in Bernardo and Smith (2000) and a non-cross-validatory version of (2) for individual models M_j , namely

$$E_{Y_{n+1}|\mathbf{Y}, M_j} \ell(Y_{n+1}, a_{M_j}(\mathbf{Y})) - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a_{M_j}(\mathbf{Y})) \xrightarrow{P} 0,$$

is established in Walker and Gutierrez-Pena (1999) where $a_{M_j}(\mathbf{Y})$ is in the action space associated with M_j .

Aside from the applications of these results to the stacking predictor, the results – if proved formally as below – establish that leave-one-out cross validation is asymptotically a Bayes optimal procedure under some conditions. It can be verified that the proofs below extend to leave- k -out cross-validation as well. That is, our results provide a Bayesian justification for using cross-validation as a way to choose a model from which to generate predictions outside of \mathcal{M} -closed problems.

For the sake of completeness, we recall that Bernardo and Smith (2000) define \mathcal{M} -closed problems as those for which a true model can be identified and

written down but is one amongst finitely many models from which an analyst has to choose. By contrast, \mathcal{M} -complete problems are those in which a true model (sometimes called a belief model) exists but is inaccessible in the sense that even though it can be conceptualized it cannot be written down or at least cannot be used directly. Effectively this means that other surrogate models must be identified and used for inferential purposes. \mathcal{M} -open problems according to Bernardo and Smith (2000) are those problems where a true model exists but cannot be specified at all.

Here however, we make a stronger distinction between \mathcal{M} -complete and \mathcal{M} -open problems by taking the view that in the \mathcal{M} -open case no true model can even be conceptualized. Hence it is inappropriate to assume the existence of a true model. We prefer this stronger distinction because it ensures that \mathcal{M} -complete and \mathcal{M} -open are disjoint classes. In both \mathcal{M} -complete and \mathcal{M} -open classes the status of the prior is unclear because none of the models under consideration are taken to be true. However, a weighting function ostensibly indistinguishable from a prior can be regarded as a sort of pseudo-belief in the sense that it is the weight one would pre-experimentally assign to the model if it were an action for predicting the outcomes of a data generator. More generally, the weights can only be interpreted as an index for a class of actions, provided the weighted combination of predictors from the J models is regarded as an action.

Although the w_j 's are often assumed to be positive and sum to one e.g., Clyde and Iversen (2013), Breiman (1996) only assumed the weights were positive and some remarks in Clyde and Iversen (2013) consider the case that the weights only satisfy a 'sum to one' constraint thereby permitting negative weights. We can see the effect of the sum to one constraint in a simple example. Following Clyde (2012), consider the two models $M_1 : Y = x_1\beta_1 + \epsilon$ and $M_2 : Y = x_2\beta_2 + \epsilon$ where the explanatory variables are orthogonal i.e., $x_1'x_2 = 0$. As shown in the Appendix at the end of the paper, if we stack these two models with the sum to one constraint we get $\hat{w}_1 = \hat{w}_2 = 1/2$. That is, predictions are generated from $Y_W = (1/2)x_1\hat{\beta}_1 + (1/2)x_2\hat{\beta}_2$ where the $\hat{\beta}_k$ are found from model M_k for $k = 1, 2$. On the other hand, if we stack M_1 and M_2 without the sum to one

constraint but with, say, a sum to two constraint we get $Y_{WO} = x_1\hat{\beta}_1 + x_2\hat{\beta}_2 + \epsilon$, i.e., $\hat{w}_1 = \hat{w}_2 = 1$. Obviously, $Y_{WO} = 2Y_W$ so Y_W is half the size it should be. This extends to three or more models and shows that the sum to one constraint can be too restrictive. In addition, permitting w_j 's to be negative increases the range of the stacking predictors and can only result in better predictions. Consequently in most of our results below we do not impose either the sum to one constraint or the non-negativity constraint.

There three main contributions of this paper are (1) a formal proof that (2) holds for several loss functions in \mathcal{M} -complete settings, (2) explicit formulae for the stacking weights for various choices of constraints on the w_j 's, and (3) a way to choose optimal basis expansions to stack so as to clarify the suggestion in Breiman (1996) that the models be chosen as different from each other as possible. In our examples, we choose two data generators, one \mathcal{M} -complete and one \mathcal{M} -open, to see how stacking performs.

The structure of this paper is as follows. In Section 2 we present the formal proof of using cross-validation to approximate posterior risk and hence derive stacking as an approximation to the Bayes action. In Section 3, we use the approximation to the posterior expected risk to derive stacking weights under several sets of constraints on the weights, observing that relaxing the non-negativity and the sum to one constraints improves prediction. In Section 4, we show how to get optimal data-driven basis expansions to stack. These bases should be different from each other in the sense of being independent; orthogonality does not seem to be helpful. In Section 5, we give two real data examples where the \mathcal{M} -complete or \mathcal{M} -open assumption is reasonable. We use them to show the effect of the sum of the coefficients and to suggest desirable properties of basis element generation. Some concluding remarks are made in Section 6.

2. Approximating posterior risk

Let $\mathcal{M} = \{M_1, \dots, M_J\}$ be a class of models and $\mathbf{y} = (y_1, \dots, y_n)$ be the vector of outcomes of $\mathbf{Y} = (Y_1, \dots, Y_n)$, where the Y_i 's are independently

distributed with probability density function (pdf) $p_j(y \mid \theta_j)$ for $j = 1 \dots, J$ equipped with a prior $w_j(\theta_j)$. Consider a loss function $\ell : \mathbb{R} \times \mathcal{A} \rightarrow \mathbb{R}$ where \mathcal{A} is the action space of a predictive decision problem. In this setting $\ell(y_{n+1}, a(\mathbf{y}))$ is the cost of taking action $a(\mathbf{y})$, where y_{n+1} is a future observation. Although the language of utility functions is more common in our context, we prefer the language of loss functions because it is more suggestive of decision theory. The posterior risk under model j is

$$\int \ell(y_{n+1}, a(\mathbf{y})) p_j(y_{n+1} \mid \mathbf{y}) dy_{n+1}, \quad (3)$$

where $p_j(\cdot \mid \mathbf{y})$ is the predictive density from model j . Given a set of convex weights $\pi(j)$ for use over the models, the overall posterior risk is

$$\int \ell(y_{n+1}, a(\mathbf{y})) p(y_{n+1} \mid \mathbf{y}) dy_{n+1} \quad (4)$$

where $p(y_{n+1} \mid \mathbf{y})$ is the predictive density marginalizing out over j as well as the θ_j 's.

The relationship between the notation in (1) and the above is that if $Y = f_j(x) + \epsilon$ we can write f_j in a generic parametric form $f_j(x) = f_j(x, \beta_j)$ so that $Y \sim p_j(y \mid \theta_j)$ means $Y \sim p_j(y \mid x, \theta_j)$ where θ_j is the concatenation of β_j and the parameters in the distribution of ϵ . We also assume without further comment that (i) the explanatory variable x and the parameter θ_j are of fixed dimension, and, for simplicity of notation, (ii) $(x, \theta_j) \in K_1 \times K_2$ where K_1 is a compact set in the space of explanatory variables and K_2 is a compact set. Strictly speaking, K_2 depends on j , but we assume that a single K_2 can be found and used for all j . This latter regularity condition can be relaxed at the cost of more notation. As a separate issue, because the Bayes predictors require integration over θ , the compactness of K_2 is only needed for the frequentist results.

In the results below we establish six versions of (2) using three different loss functions (squared error, absolute error, and logarithmic loss – also sometimes called a logarithmic scoring rule) and two different classes of predictor (Bayes and plug-in). Bayes predictors are of the form $E_j(Y_{n+1} \mid \mathbf{Y})$ and plug-in pre-

dictors are of the form $E_{\hat{\theta}_j} Y_{n+1}$ where $\hat{\theta}_j$ is an estimator of the true value of θ_j using \mathbf{Y} . To an extent the proofs of these results are similar: All of them use multiple steps of the form ‘add and subtract the right extra terms, apply the triangle inequality, and bound the result term-by-term’, and conclude by invoking a uniform integrability condition. One difference is that the results for the Bayes predictors invoke a martingale convergence theorem whereas plug-in predictors add an extra step based on the consistency of the $\hat{\theta}_j$ ’s.

We begin by giving conditions under which we can state and prove (2) for squared error and Bayes predictors; this provides formal justification for the methodology in Clyde and Iversen (2013) in \mathcal{M} -complete problems.

Theorem 2.1. *Let $\ell(z, a) = (z - a)^2$ denote squared error loss. Assume*

(i) *For any $j = 1, \dots, J$ and any pre-assigned $\epsilon > 0$,*

$$E_j(Y^{4+\epsilon}) = \int \int y^{4+\epsilon} p_j(y \mid x, \theta_j) w_j(\theta_j) d\theta_j dy < \infty,$$

and $E_j(Y^{4+\epsilon})$ is continuous for $x \in K_1$,

(ii) *For each $j = 1, \dots, J$, the conditional densities $p_j(y \mid x, \theta_j)$ are equicontinuous for $x \in K_1$ for each y and $\theta_j \in K_2$, and,*

(iii) *For each $j = 1, \dots, J$, the Bayes predictor*

$$\hat{Y}_j = E_j(Y_{n+1} \mid \mathbf{Y}) = \int \int y_{n+1} p_j(y_{n+1} \mid x_{n+1}, \theta_j) w_j(\theta_j \mid \mathbf{Y}) d\theta_j dy_{n+1}$$

is used to generate predictions at the $n + 1$ step.

Then, for any action $a(\mathbf{Y}) = \sum_{j=1}^J w_j \hat{Y}_j$, $w_j \in \mathbb{R}$ for all j , we have

$$\int \ell(y_{n+1}, a(\mathbf{Y})) p(y_{n+1} \mid \mathbf{Y}) dy_{n+1} - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y}_{-i})) \xrightarrow{L^2} 0 \text{ as } n \rightarrow \infty.$$

Proof. Fix a countable sequence $x_1, x_2, \dots \in K_1$. To establish the theorem, it is enough to show that as $n \rightarrow \infty$,

$$E_{\mathbf{Y}} \left[\int \ell(y_{n+1}, a(\mathbf{Y})) p(y_{n+1} \mid \mathbf{Y}) dy_{n+1} - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y}_{-i})) \right]^2 \rightarrow 0.$$

This can be done by adding and subtracting $(1/n) \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y}))$ to see

$$\begin{aligned} & 2E_{\mathbf{Y}} \left[E_{Y_{n+1}|\mathbf{Y}} \ell(Y_{n+1}, a(\mathbf{Y})) - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y})) \right]^2 \\ & + 2E_{\mathbf{Y}} \left[\frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y})) - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y}_{-i})) \right]^2 \end{aligned} \quad (5)$$

bounds the left hand side of the last expression from above and goes to zero.

Step 1: The first term in (5) has limit zero. To see this, recall that Walker and Gutierrez-Pena (1999) establish

$$E_{Y_{n+1}|\mathbf{Y}, M_j} \ell(Y_{n+1}, a_{M_j}(\mathbf{Y})) - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a_{M_j}(\mathbf{Y})) \xrightarrow{P} 0 \text{ as } n \rightarrow \infty,$$

for any j , where $E_{Y_{n+1}|\mathbf{Y}, M_j}$ denotes the conditional expectation with respect to $(Y_{n+1} | \mathbf{Y})$ within model M_j ; see also Clyde and Iversen (2013) for a proof valid under our hypotheses. Summing over $j = 1, \dots, J$ and recognizing that a_{M_j} is a function of \mathbf{Y} , and hence can be regarded as an element $a(\mathbf{Y})$ of the action space of \mathcal{M} , give

$$E_{Y_{n+1}|\mathbf{Y}} \ell(Y_{n+1}, a(\mathbf{Y})) - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y})) \xrightarrow{P} 0 \text{ as } n \rightarrow \infty,$$

and hence

$$\left[E_{Y_{n+1}|\mathbf{Y}} \ell(Y_{n+1}, a(\mathbf{Y})) - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y})) \right]^2 \xrightarrow{P} 0 \text{ as } n \rightarrow \infty. \quad (6)$$

Since ℓ is squared error and $a(\mathbf{Y}) = \sum_{j=1}^J w_j \hat{Y}_j$, the left hand side of (6) is

$$\begin{aligned} & \left\{ \left[E(Y_{n+1}^2 | \mathbf{Y}) - \frac{1}{n} \sum_{i=1}^n Y_i^2 \right] \right. \\ & \quad \left. + 2 \left(\sum_{j=1}^J w_j \hat{Y}_j \right) \left[\frac{1}{n} \sum_{i=1}^n Y_i - E(Y_{n+1} | \mathbf{Y}) \right] \right\}^2 \\ & = \left\{ E(Y_{n+1}^2 | \mathbf{Y}) + \left(-\frac{1}{n} \sum_{i=1}^n Y_i^2 \right) \right. \\ & \quad \left. + \left[2 \left(\sum_{j=1}^J w_j E_j(Y_{n+1} | \mathbf{Y}) \right) \left(\frac{1}{n} \sum_{i=1}^n Y_i - E(Y_{n+1} | \mathbf{Y}) \right) \right] \right\}^2, \end{aligned} \quad (7)$$

since $\hat{Y}_j = E_j(Y_{n+1} \mid \mathbf{Y})$ by Assumption (iii). Now, using $(a + b + c)^2 \leq 3(a^2 + b^2 + c^2)$ on the right hand side of (7), the left hand side of (6) is bounded by

$$\begin{aligned}
& 3E^2(Y_{n+1}^2 \mid \mathbf{Y}) + 3 \left(\frac{1}{n} \sum_{i=1}^n Y_i^2 \right)^2 \\
& + 12 \left[\sum_{j=1}^J w_j E_j(Y_{n+1} \mid \mathbf{Y}) \right]^2 \left[\frac{1}{n} \sum_{i=1}^n Y_i - E(Y_{n+1} \mid \mathbf{Y}) \right]^2 \\
& \leq 3E(Y_{n+1}^4 \mid \mathbf{Y}) + 3 \left(\frac{1}{n} \sum_{i=1}^n Y_i^2 \right)^2 \\
& + 24 \left(\sum_{j=1}^J w_j^2 \right) \left[\sum_{j=1}^J E_j^2(Y_{n+1} \mid \mathbf{Y}) \right] \\
& \quad \times \left[\left(\frac{1}{n} \sum_{i=1}^n Y_i \right)^2 + E^2(Y_{n+1} \mid \mathbf{Y}) \right]. \tag{8}
\end{aligned}$$

Using $ab \leq (a^2 + b^2)/2$ in the third term of (8) gives the new bound

$$\begin{aligned}
& 3E(Y_{n+1}^4 \mid \mathbf{Y}) + 3 \left(\frac{1}{n} \sum_{i=1}^n Y_i^2 \right)^2 \\
& + 12 \left(\sum_{j=1}^J w_j^2 \right) \left[\sum_{j=1}^J E_j^2(Y_{n+1} \mid \mathbf{Y}) \right]^2 \\
& + 12 \left(\sum_{j=1}^J w_j^2 \right) \left[\left(\frac{1}{n} \sum_{i=1}^n Y_i \right)^2 + E^2(Y_{n+1} \mid \mathbf{Y}) \right]^2. \tag{9}
\end{aligned}$$

Using $(\sum_i a_i)^2 \leq n \sum_i a_i^2$ in the third and fourth terms of (9) and then applying Cauchy-Schwarz to the results gives the upper bound

$$\begin{aligned}
& 3E(Y_{n+1}^4 \mid \mathbf{Y}) + 3 \left(\frac{1}{n} \sum_{i=1}^n Y_i^2 \right)^2 \\
& + 12J \left(\sum_{j=1}^J w_j^2 \right) \left(\sum_{j=1}^J E_j(Y_{n+1}^4 \mid \mathbf{Y}) \right) \\
& + 24 \left(\sum_{j=1}^J w_j^2 \right) \left(\frac{1}{n} \sum_{i=1}^n Y_i^2 \right)^2 + 24 \left(\sum_{j=1}^J w_j^2 \right) E(Y_{n+1}^4 \mid \mathbf{Y}). \tag{10}
\end{aligned}$$

Next, we show that each term on the right in (10) is uniformly integrable. First observe that, under Assumption (i), $E(Y_{n+1}^4 \mid \mathbf{Y})$ and $E_j(Y_{n+1}^4 \mid \mathbf{Y})$ are uniformly integrable, see Billingsley (2012), p. 498, so terms one, three, and five in (10) are uniformly integrable. Since terms two and four are nearly the same, it is enough to show term two is uniformly integrable. Begin by noting that Jensen's inequality gives

$$\left(\sum_{i=1}^n a_i \right)^{1+\epsilon} \leq n^\epsilon \left(\sum_{i=1}^n a_i^{1+\epsilon} \right), \quad (11)$$

for any $\epsilon > 0$ and $a_i \geq 0, i = 1, \dots, n$. (Set $\varphi(x) = x^{1+\epsilon}$.) Now, by Cauchy-Schwarz and (11)

$$\begin{aligned} \sup_n E \left[\left(\frac{1}{n} \sum_{i=1}^n Y_i^2 \right)^2 \right]^{1+\epsilon} &\leq \sup_n E \left(\frac{1}{n} \sum_{i=1}^n Y_i^4 \right)^{1+\epsilon} \\ &\leq \sup_n \frac{1}{n^{1+\epsilon}} E \left[n^\epsilon \left(\sum_{i=1}^n Y_i^{4(1+\epsilon)} \right) \right] \\ &\leq \sup_x E \left(Y^{4(1+\epsilon)} \right) < \infty, \end{aligned}$$

by Assumption (i) where Y denotes any Y_i as a function of x . Thus term two in (10) is uniformly integrable.

Since all the terms on the right hand side of (10) are uniformly integrable,

$$\left[E_{Y_{n+1} \mid \mathbf{Y}} \ell(Y_{n+1}, a(\mathbf{Y})) - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y})) \right]^2$$

is uniformly integrable and by (6) has limit zero in probability. Thus

$$E_{\mathbf{Y}} \left[E_{Y_{n+1} \mid \mathbf{Y}} \ell(Y_{n+1}, a(\mathbf{Y})) - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y})) \right]^2 \rightarrow 0 \quad (12)$$

as $n \rightarrow \infty$ by the Theorem 25.12 in Billingsley (2012). So the first term in (5) also goes to 0 as $n \rightarrow \infty$.

Step 2: Obtain a bound on the second term in (5). Using $(\sum_i a_i)^2 \leq n \sum_i a_i^2$ it is seen that the second term in (5) is bounded by twice

$$\frac{1}{n} E_{\mathbf{Y}} \sum_{i=1}^n [\ell(Y_i, a(\mathbf{Y})) - \ell(Y_i, a(\mathbf{Y}_{-i}))]^2. \quad (13)$$

The terms in (13) depend on $x_1, \dots, x_n \in K$ and the i -th term depends on x_i differently from how it depends on the $x_{i'}$'s for $i' \neq i$. However, the sum is symmetric in the x_i 's.

Recalling $\ell(z, a) = (z - a)^2$ and $a(\mathbf{Y}) = \sum_{j=1}^J w_j \hat{Y}_j$, where w_j is a coefficient assigned to model M_j , the i -th term in (13) is

$$E_{\mathbf{Y}} \left\{ \left[\left(\sum_{j=1}^J w_j \hat{Y}_j \right)^2 - \left(\sum_{j=1}^J w_j \hat{Y}_{j,-i} \right)^2 \right] + 2Y_i \left(\sum_{j=1}^J w_j \hat{Y}_{j,-i} - \sum_{j=1}^J w_j \hat{Y}_j \right) \right\}^2. \quad (14)$$

Using $(a + b)^2 \leq 2a^2 + 2b^2$, $(a^2 - b^2) = (a - b)(a + b)$ and Cauchy-Schwarz, (14) is bounded from above by

$$2E_{\mathbf{Y}} \left[\left(\sum_{j=1}^J w_j \hat{Y}_j - \sum_{j=1}^J w_j \hat{Y}_{j,-i} \right)^2 \left(\sum_{j=1}^J w_j \hat{Y}_j + \sum_{j=1}^J w_j \hat{Y}_{j,-i} \right)^2 \right] + 8\sqrt{E_{\mathbf{Y}} Y_i^4} \left[E_{\mathbf{Y}} \left(\sum_{j=1}^J w_j (\hat{Y}_j - \hat{Y}_{j,-i}) \right)^4 \right]^{1/2}. \quad (15)$$

Re-arranging in (15) and using Cauchy-Schwarz and Jensen's inequality repeatedly gives the upper bound

$$\begin{aligned} & 2E_{\mathbf{Y}} \left[\sum_{j=1}^J w_j (\hat{Y}_j - \hat{Y}_{j,-i}) \right]^2 \left[\sum_{j=1}^J w_j (\hat{Y}_j + \hat{Y}_{j,-i}) \right]^2 \\ & + 8\sqrt{E_{\mathbf{Y}} Y_i^4} \left[E_{\mathbf{Y}} J^2 \left(\sum_{j=1}^J w_j^2 (\hat{Y}_j - \hat{Y}_{j,-i})^2 \right)^2 \right]^{1/2} \\ & \leq 2J \left(\sum_{j=1}^J w_j^2 \right)^2 \left[E_{\mathbf{Y}} \sum_{j=1}^J (\hat{Y}_j + \hat{Y}_{j,-i})^4 \right]^{1/2} \left[E_{\mathbf{Y}} \sum_{j=1}^J (\hat{Y}_j - \hat{Y}_{j,-i})^4 \right]^{1/2} \\ & + 8J\sqrt{E_{\mathbf{Y}} Y_i^4} \left(\sum_{j=1}^J w_j^4 \right)^{1/2} \left[E_{\mathbf{Y}} \sum_{j=1}^J (\hat{Y}_j + \hat{Y}_{j,-i})^4 \right]^{1/2}. \end{aligned}$$

This last expression equals

$$\left[2J \left(\sum_{j=1}^J w_j^2 \right)^2 \left(E_{\mathbf{Y}} \sum_{j=1}^J \left(\hat{Y}_j + \hat{Y}_{j,-i} \right)^4 \right)^{1/2} + 8J \sqrt{E_{\mathbf{Y}} Y_i^4} \left(\sum_{j=1}^J w_j^4 \right)^{1/2} \right] \times \left[E_{\mathbf{Y}} \sum_{j=1}^J \left(\hat{Y}_j - \hat{Y}_{j,-i} \right)^4 \right]^{1/2} \quad (16)$$

where $\hat{Y}_j = E_j(Y_{n+1} \mid \mathbf{Y})$ and $\hat{Y}_{j,-i} = E_j(Y_{n+1} \mid \mathbf{Y}_{-i})$ by Assumption (iii). Next, we show the second factor in (16) goes to zero as $n \rightarrow \infty$.

Step 3, Part 1: The fourth power inside the expectation in (16) goes to zero almost everywhere. First observe that there is a σ -field σ_∞ so that $\sigma(Y_1, \dots, Y_n) \nearrow \sigma_\infty$ as $n \rightarrow \infty$. Similarly, $\sigma(Y_1, \dots, Y_{n-1}) \nearrow \sigma_\infty$. Indeed, the σ -field generated by any $n-1$ of the Y_{-i} 's converges to σ_∞ for each j . Since $E_j Y_{n+1}$ is finite for all j , the martingale convergence theorem, Billingsley (2012) (Theorem 35.6 on p. 499), gives that, for each j ,

$$E_j(Y_{n+1} \mid \mathbf{Y}), E_j(Y_{n+1} \mid \mathbf{Y}_{-i}) \rightarrow E_j(Y_{n+1} \mid \sigma_\infty)$$

almost everywhere in M_j as $n \rightarrow \infty$. Hence,

$$E_j(Y_{n+1} \mid \mathbf{Y}) - E_j(Y_{n+1} \mid \mathbf{Y}_{-i}) \rightarrow 0, \quad (17)$$

almost everywhere in M_j for each of the x_i 's. However, since all the models have the same sets of measure zero in the underlying measure space, convergence (17) holds almost everywhere in the limit over the mixture of all J models, i.e., with respect to $p(\mathbf{y})$ in a limiting sense. Thus, the fourth powers also converge to zero almost everywhere in the mixture of the J models.

Step 3, Part 2: The fourth powers inside the expectation in (16) are uni-

formly integrable. By the Cauchy-Schwarz inequality,

$$\begin{aligned}
& \sum_{j=1}^J (E_j(Y_{n+1} \mid \mathbf{Y}) - E_j(Y_{n+1} \mid \mathbf{Y}_{-i}))^4 \\
& \leq 4 \sum_{j=1}^J (E_j(Y_{n+1}^2 \mid \mathbf{Y}) + E_j(Y_{n+1}^2 \mid \mathbf{Y}_{-i}))^2 \\
& \leq 8 \sum_{j=1}^J (E_j(Y_{n+1}^4 \mid \mathbf{Y}) + E_j(Y_{n+1}^4 \mid \mathbf{Y}_{-i})). \tag{18}
\end{aligned}$$

So, by Assumption (i) and Billingsley (2012) (Lemma p. 498), we have that the right hand side is a uniformly integrable sequence since the sequence of x_i 's in K_1 is regarded as a countable collection of fixed design points. Assumption (ii) together with Step 3, Part 1, gives

$$E_{\mathbf{Y}} \left[\sum_{j=1}^J (E_j(Y_{n+1} \mid \mathbf{Y}) - E_j(Y_{n+1} \mid \mathbf{Y}_{-i}))^4 \right] \rightarrow 0 \text{ as } n \rightarrow \infty \tag{19}$$

uniformly in the x_i 's. Thus, the bound on the i -th term in (13) is independent of i and so is a valid upper bound that goes to zero for all n terms in (13), therefore bounding the average.

To conclude the proof, we show that the first factor in (16) is uniformly bounded as a function of n for x_1, \dots, x_n in a compact set. Indeed, by Assumption (i), the expectation $E_{\mathbf{Y}} Y_i^4$ is

$$\int y^4 \left(\sum_{j=1}^J \pi_j \int p_j(y \mid x_i, \theta_j) w_j(\theta_j) d\theta_j \right) dy = \sum_{j=1}^J \pi_j E_j(Y^4) < \infty. \tag{20}$$

Moreover, $\sum_{j=1}^J (\hat{Y}_j + \hat{Y}_{j,-i})^4$ is uniformly integrable for the same reason as (18) is and converges to

$$16 \sum_{j=1}^J [E_j(Y_{n+1} \mid \sigma_{\infty})]^4.$$

So, by Billingsley (2012) (Theorem 25.12 p. 361) $16 \sum_{j=1}^J [E_j(Y_{n+1} \mid \sigma_{\infty})]^4$ is integrable and $E_{\mathbf{Y}} \sum_{j=1}^J (\hat{Y}_j + \hat{Y}_{j,-i})^4 \rightarrow E_{\mathbf{Y}} 16 \sum_{j=1}^J [E_j(Y_{n+1} \mid \sigma_{\infty})]^4$ giving that the first factor in (16) is bounded.

Taken together, these statements imply that the second term in (5) goes to 0 as $n \rightarrow \infty$. □

The result in Theorem 2.1 remains true for squared error loss if we use plug-in predictors of the form

$$\hat{Y}_j = E_{\hat{\theta}_j(\mathbf{Y})}(Y_{n+1}) = \int y_{n+1} p_{\hat{\theta}_j(\mathbf{Y})}(y_{n+1}) dy_{n+1},$$

where $\hat{\theta}_j$ is any consistent estimator for θ_j in M_j rather than Bayes predictors as in Assumption (iii). This assertion is in the following.

Theorem 2.2. *Let $\ell(z, a) = (z - a)^2$ denote squared error loss. Assume*

(i) *For any $j = 1, \dots, J$ and any pre-assigned $\epsilon > 0$,*

$$E_j(Y^{4+\epsilon}) = \int \int y^{4+\epsilon} p_j(y \mid x, \theta_j) w_j(\theta_j) d\theta_j dy < \infty,$$

and $E_j(Y^{4+\epsilon})$ is continuous for $x \in K_1$,

(ii) *For each $j = 1, \dots, J$, the conditional densities $p_j(y \mid x, \theta_j)$ are equicontinuous for $x \in K_1$ for each y and $\theta_j \in K_2$, and,*

(iii) *For each $j = 1, \dots, J$, let the plug-in predictor*

$$\hat{Y}_j = E_{\hat{\theta}_j(\mathbf{Y})}(Y_{n+1}) = \int y_{n+1} p_{\hat{\theta}_j(\mathbf{Y})}(y_{n+1}) dy_{n+1},$$

be used to generate predictions at the $n+1$ step, where $\hat{\theta}_j(\mathbf{Y})$ is a consistent estimator for θ_j , and

(iv) *For each $j = 1, \dots, J$, $E_{\theta_j}(Y_{n+1})$ and $E_{\theta_j}(Y_{n+1}^4)$ are continuous for $\theta_j \in \Theta_j$, where $\Theta_j \subset K_2$ is a compact parameter space for M_j .*

Then, for any action $a(\mathbf{Y}) = \sum_{j=1}^J w_j \hat{Y}_j$, $w_j \in \mathbb{R}$ for all j , we have

$$\int \ell(y_{n+1}, a(\mathbf{Y})) p(y_{n+1} \mid \mathbf{Y}) dy_{n+1} - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y}_{-i})) \xrightarrow{L^2} 0 \text{ as } n \rightarrow \infty.$$

Proof. Fix a countable sequence $x_1, x_2, \dots \in K_1$. Recall expression (5). The first term is treated much the same as in Theorem 2.1. The main difference is that Assumption (iv) can be used to establish uniform integrability of $E_{\hat{\theta}_j(\mathbf{Y})} Y_{n+1}^4$ because it gives $\sup_n E_{\mathbf{Y}} \left[E_{\hat{\theta}_j(\mathbf{Y})} (Y_{n+1}^4) \right]^2 < \infty$.

Showing the second term in (5) goes to zero uses the plug-in estimator version of (16). Showing it goes to zero requires

$$E_{\hat{\theta}_j(\mathbf{Y})}(Y_{n+1}) \xrightarrow{L_1} E_{\theta_j}(Y_{n+1}) \text{ as } n \rightarrow \infty \quad (21)$$

which uses Assumption (iv) for the first moment (rather than a martingale argument) and a uniform integrability argument similar to that used in Step 3 of the proof of Theorem 2.1. □

Unsurprisingly, the conclusions of Theorems 2.1 and 2.2 continue to hold if the squared error is replaced by the absolute error $\ell(z, a) = |z - a|$. We state the Bayes and plug-in versions for absolute error in the following.

Theorem 2.3. *Let $\ell(z, a) = |z - a|$ denote absolute error loss. Assume*

(i) *For any $j = 1, \dots, J$ and any pre-assigned $\epsilon > 0$,*

$$E_j(Y^{2+\epsilon}) = \int \int y^{2+\epsilon} p_j(y \mid x, \theta_j) w_j(\theta_j) d\theta_j dy < \infty,$$

and $E_j(Y^{2+\epsilon})$ is continuous for $x \in K_1$,

(ii) *For each $j = 1, \dots, J$, the conditional densities $p_j(y \mid x, \theta_j)$ are equicontinuous for $x \in K_1$ for each y and $\theta_j \in K_2$, and,*

(iii) *For each $j = 1, \dots, J$, let either the Bayes or the plug-in predictor from Theorem 2.1 or Theorem 2.2, respectively, be used to generate predictions at the $n + 1$ time step.*

(iv) *If plug-in predictors are chosen in Assumption (iii), then assume in addition that for each $j = 1, \dots, J$, $E_{\theta_j}(Y_{n+1})$ and $E_{\theta_j}(Y_{n+1}^2)$ are continuous as functions of $\theta_j \in \Theta_j$, where $\Theta_j \subset K_2$ is a compact parameter space for M_j .*

Then, for any action $a(\mathbf{Y}) = \sum_{j=1}^J w_j \hat{Y}_j$, $w_j \in \mathbb{R}$ for all j , we have

$$\int \ell(y_{n+1}, a(\mathbf{Y})) p(y_{n+1} | \mathbf{Y}) dy_{n+1} - \frac{1}{n} \sum_{i=1}^n \ell(Y_i, a(\mathbf{Y}_{-i})) \xrightarrow{L^2} 0 \text{ as } n \rightarrow \infty.$$

Proof. Fix a countable sequence $x_1, x_2, \dots \in K_1$. As before, we recall (5). Verifying that the first term goes to zero is the same as in the proof of Theorem 2.1 for the Bayes posterior mean predictor and can be modified as explained in the proof of Theorem 2.2 for the plug-in predictor.

The i -th summand in the second term in (5) is bounded by

$$E_{\mathbf{Y}} \left(\left| Y_i - \sum_{j=1}^J w_j \hat{Y}_j \right| - \left| Y_i - \sum_{j=1}^J w_j \hat{Y}_{j,-i} \right| \right)^2. \quad (22)$$

Since $(|a| - |b|)^2 \leq (a - b)^2$, (22) is bounded from above by

$$E_{\mathbf{Y}} \left[\sum_{j=1}^J w_j \left(\hat{Y}_j - \hat{Y}_{j,-i} \right) \right]^2. \quad (23)$$

Using Cauchy-Schwarz inequality, (23) is upper bounded by

$$E_{\mathbf{Y}} \left\{ \left(\sum_{j=1}^J w_j^2 \right) \left[\sum_{j=1}^J \left(\hat{Y}_j - \hat{Y}_{j,-i} \right)^2 \right] \right\} = \left(\sum_{j=1}^J w_j^2 \right) \sum_{j=1}^J E_{\mathbf{Y}} \left(\hat{Y}_j - \hat{Y}_{j,-i} \right)^2.$$

If the Bayes predictors are used, the last expression can be controlled by Step 3 in the proof of Theorem 2.1. If the plug-in predictors are used, then the last expression can be controlled by the extra step mentioned in the proof of Theorem 2.2, see (21). □

The so-called log-loss is qualitatively different from squared error or absolute error because log-loss can be positive or negative. It is therefore better regarded as a utility function even though in some cases it can be physically interpreted as the code length function for a Shannon code. Here we extend our results to the log-utility to verify that cross-validation continues to remain an asymptotically Bayes procedure. Now, an action a is of the form

$$a(Y_{n+1} | \mathbf{Y}) = \sum_{j=1}^J w_j p_j(Y_{n+1} | \mathbf{Y}), \quad (24)$$

and the corresponding log-utility is

$$u(Y_{n+1}, a(Y_{n+1} | \mathbf{Y})) = \log \left[\sum_{j=1}^J w_j p_j(Y_{n+1} | \mathbf{Y}) \right]. \quad (25)$$

As shown in our next result, cross-validation approximates the posterior expected utility of the Bayes action of the form (24) or the plug-in action of the form $a(Y_{n+1} | \mathbf{Y}) = \sum_{j=1}^J w_j p_{\hat{\theta}_j(\mathbf{Y})}(Y_{n+1})$ where $\hat{\theta}_j(\mathbf{Y})$ is a consistent estimator of $\theta_j \in K_2$ in M_j .

Theorem 2.4. *Let $u(Y_{n+1}, a(\mathbf{Y})) = \log \left[\sum_{j=1}^J w_j p_j(Y_{n+1} | \mathbf{Y}) \right]$ denote log-utility. Assume*

(i) *For each $j = 1, \dots, J$, there is a function $B_j(\cdot)$ so that*

$$\sup_{\mathbf{Y}} |\log p_j(Y_{n+1} | \mathbf{Y})| \leq B_j(Y_{n+1}) < \infty,$$

$B_j(\cdot)$ is independent of x_1, x_2, \dots , and

$$E[g(Y_{n+1})] < \infty,$$

where

$$g(Y_{n+1}) = \max \left\{ \left(\log \sum_{j=1}^J w_j e^{-B_j(Y_{n+1})} \right)^4, \left(\log \sum_{j=1}^J w_j e^{B_j(Y_{n+1})} \right)^4 \right\}.$$

(ii) *For each $j = 1, \dots, J$, the conditional densities $p_j(y | x, \theta_j)$ are equicontinuous in x for each y and $\theta_j \in \Theta_j \subset K_2$, and the predictive densities $p_j(y | \mathbf{Y})$ within the j -th model are uniformly equicontinuous in y .*

(iii) *For each $j = 1, \dots, J$, let the Bayes action (24) be used to generate predictions at the $n + 1$ time step.*

Then, we have

$$\int u(y_{n+1}, a(\mathbf{Y})) p(y_{n+1} | \mathbf{Y}) dy_{n+1} - \frac{1}{n} \sum_{i=1}^n u(Y_i, a(\mathbf{Y}_{-i})) \xrightarrow{L^2} 0 \text{ as } n \rightarrow \infty.$$

If $p_{\hat{\theta}_j(\mathbf{Y})}(y_{n+1})$ where $\hat{\theta}_j(\mathbf{Y})$ is a consistent estimator of θ_j is used instead of $p_j(y_{n+1} | \mathbf{Y})$, the result still holds.

Proof. The structure of the proof is similar to that of Theorem 2.1 or Theorem 2.2. For both Bayes and plug-in predictors the main difference between the proof of this theorem and the proofs of Theorems 2.1 and 2.2 is that in *Step 1* Assumption (i) is used to establish uniform integrability (and convergence to zero) of the first term in (5). Showing that the second term in (5) goes to zero for the Bayes predictors requires Assumptions (i) and (ii) to get

$$p_j(y_{n+1} \mid \mathbf{Y}) \rightarrow p_j(y_{n+1} \mid \sigma_\infty)$$

almost everywhere as $n \rightarrow \infty$ to set up an application of the dominated convergence theorem. For the plug-in predictors, we need the extra step described in the proof of Theorem 2.2, see (21). \square

3. Derivation of stacking weights

Suppose we have J predictors $\hat{Y}_1, \dots, \hat{Y}_J$ from distinct models. Then we might seek weights \hat{w}_j , using the training data, so as to form a model average prediction at x_{new} of the form

$$\hat{y}(x_{new}) = \sum_{j=1}^J \hat{w}_j \hat{y}_j(x_{new}). \quad (26)$$

From a Bayesian point of view, one should find the action that minimizes the posterior risk (or maximizes the posterior expected utility) given the data \mathbf{y} . Theorem 2.1 shows that the posterior risk is asymptotically equivalent to

$$\frac{1}{n} \sum_{i=1}^n \ell(y_i, a(\mathbf{y}_{-i})) = \frac{1}{n} \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i}(x_i) \right)^2,$$

when ℓ is squared error loss. Ignoring the $(1/n)$ and minimizing over the \hat{w}_j 's gives the same expression as (1). That is, the stacking weights are asymptotically Bayes optimal – the precise form of optimality given by the constraints imposed on the w_j 's – and can be used in (26) to give the stacking predictor. This formalizes the heuristic approximations used in Clyde and Iversen (2013).

There are two constraints on the w_j 's that are commonly used. One is the ‘sum to one’ constraint that requires $\sum_j w_j = 1$ (see Clyde and Iversen

(2013)) and the other is the non-negativity constraint that requires all $w_j \geq 0$ (see Breiman (1996)). Removing the non-negativity constraint and relaxing the sum-to-one constraint to a sum-to- m constraint give the following.

Theorem 3.1. *The weights w_1, \dots, w_J achieving*

$$\min_w \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i}(x_i) \right)^2 \text{ subject to } \sum_{j=1}^J w_j = m$$

are of the form

$$\hat{w} \propto U^{-1} 1_J,$$

where

$$\begin{aligned} U &= (u_{lj})_{J \times J}, \\ u_{lj} &= \sum_{i=1}^n \left(\frac{y_i}{m} - \hat{y}_{j,-i} \right) \hat{y}_{l,-i} - \sum_{i=1}^n (y_i - \hat{y}_{j,-i}) y_i, \\ 1_J &= (1, \dots, 1)'. \end{aligned} \tag{27}$$

Proof. This is a standard Lagrange multipliers problem. Write the Lagrangian as

$$L = - \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i} \right)^2 - \lambda_0 \left(\sum_{j=1}^J w_j - m \right).$$

Then \hat{w} is the solution of the following system,

$$\begin{aligned} \frac{\partial L}{\partial w_l} &= 2 \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i} \right) \hat{y}_{l,-i} - \lambda_0 = 0 \text{ for } l = 1, \dots, J, \\ \frac{\partial L}{\partial \lambda_0} &= \sum_{j=1}^J w_j - m = 0. \end{aligned} \tag{28}$$

From (28) and (29), we have

$$\begin{aligned}
& \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i} \right) \hat{y}_{l,-i} = \frac{\lambda_0}{2} \\
& \Rightarrow \sum_{i=1}^n y_i \hat{y}_{j,-i} - \sum_{j=1}^J w_j \sum_{i=1}^n \hat{y}_{j,-i} \hat{y}_{l,-i} = \frac{\lambda_0}{2} \\
& \Rightarrow \frac{1}{m} \sum_{i=1}^n y_i \hat{y}_{j,-i} \sum_{j=1}^J w_j - \sum_{j=1}^J w_j \sum_{i=1}^n \hat{y}_{j,-i} \hat{y}_{l,-i} - \sum_{j=1}^J \sum_{i=1}^n (y_i - \hat{y}_{j,-i}) y_i w_j \\
& \quad = \frac{\lambda_0}{2} - \sum_{j=1}^J \sum_{i=1}^n (y_i - \hat{y}_{j,-i}) y_i w_j \quad \text{for } l = 1, \dots, J
\end{aligned}$$

Since the right hand side does not depend on l , we have

$$\frac{1}{m} \sum_{i=1}^n y_i \hat{y}_{j,-i} \sum_{j=1}^J w_j - \sum_{j=1}^J w_j \sum_{i=1}^n \hat{y}_{j,-i} \hat{y}_{l,-i} - \sum_{j=1}^J \sum_{i=1}^n (y_i - \hat{y}_{j,-i}) y_i w_j \propto 1.$$

Rearranging gives

$$\begin{aligned}
& w_1 \left(\frac{1}{m} \sum_{i=1}^n y_i \hat{y}_{l,-i} - \sum_{i=1}^n \hat{y}_{1,-i} \hat{y}_{l,-i} - \sum_{i=1}^n (y_i - \hat{y}_{1,-i}) y_i \right) \\
& + w_2 \left(\frac{1}{m} \sum_{i=1}^n y_i \hat{y}_{l,-i} - \sum_{i=1}^n \hat{y}_{2,-i} \hat{y}_{l,-i} - \sum_{i=1}^n (y_i - \hat{y}_{2,-i}) y_i \right) \\
& \vdots \\
& + w_J \left(\frac{1}{m} \sum_{i=1}^n y_i \hat{y}_{l,-i} - \sum_{i=1}^n \hat{y}_{J,-i} \hat{y}_{l,-i} - \sum_{i=1}^n (y_i - \hat{y}_{J,-i}) y_i \right) \propto 1,
\end{aligned}$$

for $l = 1, \dots, J$.

In matrix form, this system of equations is

$$Uw \propto 1_J,$$

where U and 1_J are defined as in (27). Therefore, the solution is

$$\hat{w} \propto U^{-1} 1_J,$$

which can be rescaled to satisfy the sum to m constraint.

□

Corollary 3.1. *If $m = 1$, then the weights w_1, \dots, w_J achieving*

$$\min_w \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i}(x_i) \right)^2 \quad \text{subject to} \quad \sum_{j=1}^J w_j = 1$$

are of the form

$$\hat{w} \propto (\hat{e}'\hat{e})^{-1} 1_J, \quad (30)$$

where

$$\hat{e} = (y_i - \hat{y}_{j,-i})_{n \times J} \quad \text{and} \quad 1_J = (1, \dots, 1)'.$$

Remark 3.1. *This corollary is the result Clyde and Iversen (2013) used.*

For contrast, let us solve (1) but without any sum constraint (and without the non-negativity constraint). Now, the Lagrangian is

$$L = - \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i} \right)^2$$

and \hat{w} is the solution of the system of equations

$$\frac{\partial L}{\partial w_l} = 2 \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i} \right) \hat{y}_{l,-i} = 0 \quad \text{for } l = 1, \dots, J.$$

Therefore,

$$\begin{aligned} \sum_{i=1}^n \hat{y}_{l,-i} \sum_{j=1}^J w_j \hat{y}_{j,-i} &= \sum_{i=1}^n y_i \hat{y}_{l,-i} \\ \Leftrightarrow \sum_{j=1}^J \left(\sum_{i=1}^n \hat{y}_{l,-i} \hat{y}_{j,-i} \right) w_j &= \sum_{i=1}^n y_i \hat{y}_{l,-i}, \quad \text{for } l = 1, \dots, J, \end{aligned} \quad (31)$$

or, in matrix form,

$$Tw = c,$$

where

$$\begin{aligned} T &= \left(\sum_{i=1}^n \hat{y}_{l,-i} \hat{y}_{j,-i} \right)_{J \times J}, \\ c &= \left(\sum_{i=1}^n y_i \hat{y}_{1,-i}, \dots, \sum_{i=1}^n y_i \hat{y}_{J,-i} \right)'. \end{aligned} \quad (32)$$

Hence the solution to (1) without the sum to one constraint and without the non-negativity constraint is

$$\hat{w} = T^{-1}c.$$

We summarize this in the following theorem.

Theorem 3.2. *The weights w_1, \dots, w_J achieving*

$$\min_w \sum_{i=1}^n \left(y_i - \sum_{j=1}^J w_j \hat{y}_{j,-i}(x_i) \right)^2$$

are of the form

$$\hat{w} = T^{-1}c, \tag{33}$$

where T and c are given in (32). In addition, if the J predictors are orthonormal,

$$\sum_{i=1}^n \hat{y}_{l,-i} \hat{y}_{j,-i} = \delta_{l \neq j} \quad (1 \text{ if } l \neq j \text{ and } 0 \text{ otherwise}),$$

then $T = I$ and the solution becomes

$$\hat{w}_j = \sum_{i=1}^n y_i \hat{y}_{j,-i}, \text{ for } j = 1, \dots, J. \tag{34}$$

Note that the minimum in Corollary 3.1 with the sum to one constraint is taken over a smaller set than that of Theorem 3.2 without any sum constraint. So, when the stacking weights from the two cases both exist, we expect the latter to give better predictive performance because the minimum in Theorem 3.2 can only be smaller than the minimum in Corollary 3.1. Hence we do not favor imposing the sum to one constraint. Indeed, we find in our computed examples that when a sum to one constraint gives better prediction, it is merely a happenstance from the more general optimization. This is straightforward because if we find the optimal weights from Theorem 3.2 then we can use them to find $m = \sum_{j=1}^J w_j$ for use in Theorem 3.1.

Using arguments similar to those used in the proof of Theorem 3.2, the following result extends Theorem 3.2 to a Hilbert space \mathcal{H} equipped with an

empirical inner product

$$\langle g, h \rangle_n = \frac{1}{n} \sum_{i=1}^n g(x_i)h(x_i) \quad \forall g, h \in \mathcal{H}.$$

Theorem 3.3. *The weights w_1, \dots, w_J achieving*

$$\min_w \sum_{i=1}^n \left(y(x_i) - \sum_{j=1}^J w_j \hat{f}_{j,-i}(x_i) \right)^2,$$

where y and $\hat{f}_{j,-i}$, $j = 1, \dots, J$, belong to \mathcal{H} , are of the form

$$\hat{w} = T^{-1}c,$$

where T and c are of the same form as (32).

As $n \rightarrow \infty$, there are conditions that ensure the empirical inner product $\langle g, h \rangle_n$ converges uniformly to the inner product $\langle g, h \rangle = \int g(x)h(x)dx$ of the \mathcal{H} space, see van de Geer (2014). Therefore, as n increases we can approximate the empirical inner product by the \mathcal{H} inner product and the results in Theorem 3.3 will remain true.

4. What models should we put in the stack?

Here, we show that the intuition of Breiman (1996) that the models to be stacked should be as different as possible is only partially correct. What matters about the models to be stacked is that they be independent. The extra ‘difference’ amongst models from imposing orthogonality is not actually helpful in terms of reducing the error criterion (1). We show this for models constructed in general Hilbert spaces of functions and then provide one possible answer for how to construct the models in a Hilbert space to be stacked. Note that this restriction limits us to \mathcal{M} -complete problems. In general, in \mathcal{M} -open problems one cannot assume the regression function is in a Hilbert space. However, in our \mathcal{M} -open example in Subsec. 5.2 we did not find orthonormality of a basis gave better predictions.

4.1. The error depends only on the span of the model list.

In the last section, we saw that releasing the sum to one constraint can only reduce the error criterion and from our example in the introduction we saw that this constraint can often be genuinely harmful. This argument is particularly strong outside \mathcal{M} -closed settings where model mis-specification is always present.

Our first result shows that given a set of models to stack, the error depends only on the span of the models; requiring that the models to be stacked be orthogonal as well as independent does not reduce the error. Our result is the following.

Theorem 4.1. *Let \mathcal{H} be a Hilbert space with inner product denoted $\langle \cdot, \cdot \rangle$. Let $\mathcal{M} = \{f_1, \dots, f_J\}$ and $\mathcal{M}' = \{f'_1, \dots, f'_{J'}\}$ be sets of elements from \mathcal{H} with minima $Q_{\min}^{\mathcal{M}}$ and $Q_{\min}^{\mathcal{M}'}$ for (1), respectively. Denote the span of a set of elements in \mathcal{H} by $\langle \cdot \rangle$. Then, if $\langle \mathcal{M} \rangle = \langle \mathcal{M}' \rangle$,*

$$Q_{\min}^{\mathcal{M}} = Q_{\min}^{\mathcal{M}'},$$

i.e., the stacking error only depends on the span of the predictors.

Proof. This involves routine manipulations with Hilbert spaces, see the Appendix for details. \square

Theorem 4.1 means that given a fixed subspace $S \subset \mathcal{H}$, any basis for S is as good as any other for forming a stacking predictor. So, we are free to choose whichever basis is most convenient.

Note that Theorem 4.1 only applies in the absence of constraints on the coefficients w_j . Indeed, the conclusion may be false if constraints are imposed. Let $J = J'$, $\mathcal{M} = \{\hat{y}_j = (\hat{y}_{j,-1}, \dots, \hat{y}_{j,-n})', j = 1, \dots, J\}$ be an orthogonal basis, and $\mathcal{M}' = \{\hat{y}'_j = (\hat{y}'_{j,-1}, \dots, \hat{y}'_{j,-n})', j = 1, \dots, J\}$ be any basis of $\langle \mathcal{M}' \rangle =$

$\langle \mathcal{M} \rangle$. Then, under the sum to one constraint on \mathcal{M} we have

$$\begin{aligned}
Q_{\min}^{\mathcal{M}} &= \left\| y - \sum_{j=1}^J \hat{w}_j \hat{y}_j \right\|^2 \\
&= \left\| \sum_{j=1}^J \langle y, \hat{y}_j \rangle \hat{y}_j + \sum_{j=J+1}^n \langle y, e_j \rangle e_j - \sum_{j=1}^J \hat{w}_j \hat{y}_j \right\|^2 \\
&= \left\| \sum_{j=1}^J (\langle y, \hat{y}_j \rangle - \hat{w}_j) \hat{y}_j \right\|^2 + \|y_2\|^2,
\end{aligned} \tag{35}$$

where \hat{w} is now the solution in Corollary 3.1, $\{e_j, j = J+1, \dots, n\}$ are complement vectors of $\{\hat{y}_j, j = 1, \dots, J\}$ to form an orthonormal basis of \mathbb{R}^n , and $y = y_1 + y_2 = \sum_{j=1}^J \langle y, \hat{y}_j \rangle \hat{y}_j + \sum_{j=J+1}^n \langle y, e_j \rangle e_j$. Similarly, for \mathcal{M}' we have

$$Q_{\min}^{\mathcal{M}'} = \left\| \sum_{j=1}^J (\alpha_j - \hat{w}'_j) \hat{y}'_j \right\|^2 + \|y_2\|^2, \tag{36}$$

where \hat{w}' is the solution in Corollary 3.1 and $y = y_1 + y_2 = \sum_{j=1}^J \alpha_j \hat{y}'_j + \sum_{j=J+1}^n \langle y, e_j \rangle e_j$. Obviously, from (35) and (36), it is possible for $Q_{\min}^{\mathcal{M}} < Q_{\min}^{\mathcal{M}'}$ or $Q_{\min}^{\mathcal{M}} > Q_{\min}^{\mathcal{M}'}$. This can be seen from the following example. Let $J = J' = 1$ and $\hat{y}' = k\hat{y}$, then $\hat{w} = \hat{w}' = 1$ and $\alpha = \langle y, \hat{y} \rangle / k$. Hence $Q_{\min}^{\mathcal{M}} = (\langle y, \hat{y} \rangle - 1)^2 + \|y_2\|^2$ and $Q_{\min}^{\mathcal{M}'} = (\langle y, \hat{y} \rangle - k)^2 + \|y_2\|^2$. So, by careful choice of k , $Q_{\min}^{\mathcal{M}}$ can be larger than $Q_{\min}^{\mathcal{M}'}$ or the reverse.

To reinforce Theorem 4.1, we observe that reducing the dimension of the span of the predictors can only increase the error criterion.

Theorem 4.2. *Let $\mathcal{M} = \{f_1, \dots, f_J\}$ be a basis and $\mathcal{N} = \{f_1, \dots, f_{J-1}\}$. Let $Q_{\min}^{\mathcal{M}}$ and $Q_{\min}^{\mathcal{N}}$ be the minima of (1) corresponding to \mathcal{M} and \mathcal{N} , respectively. Then,*

$$Q_{\min}^{\mathcal{M}} \leq Q_{\min}^{\mathcal{N}}. \tag{37}$$

Proof. This involves relatively routine manipulations with Hilbert spaces, see the Appendix for details. \square

Taken together, Theorem 4.1 and Theorem 4.2 tell us that the predictors being stacked should be different from one another in the sense of being independent (but not necessarily orthogonal) and that the stacking error (1) is a non-increasing function of the span of the predictors. Thus, when choosing predictors to stack, there is a tradeoff between the number of predictors and their proximity to a true model assuming one exists. That is, using more predictors will generally be helpful, but using fewer, better predictors can easily outperform many, weaker predictors.

4.2. Optimal choice of predictors to stack.

Having seen that both the number of basis elements and the proximity of a linear combination of them to a true function (if a true function exists) can be important we want to choose the basis elements effectively. The results in Subsec. 4.1 mean that, without loss of generality, we can limit our search to orthogonal bases. Hence, in this subsection, we propose a data-driven method to choose an optimal number of basis elements even if the set of basis elements is not unique.

Assume we have an orthonormal basis for a space $\langle \{e_1, \dots, e_J\} \rangle$ then for each $J' \leq J$ we can form

$$\hat{y}_{J', \sigma_k, \lambda, i}(x_{\sigma_k(i)}) = \sum_{j=1}^{J'} \langle e_j, \hat{f}_\lambda(\cdot \mid x_{\sigma_k(1)}, \dots, x_{\sigma_k(i-1)}) \rangle e_j(x_{\sigma_k(i)}), \quad (38)$$

where \hat{f}_λ is an estimate of the true predictor, for instance from the Nadaraya-Watson nonparametric regression estimator, see Nadaraya (1964) and Watson (1964), λ is a tuning parameter, and σ_k for $k = 1, \dots, K$ is a collection of independent permutations of $\{1, \dots, n\}$. Then, in principle, we can find

$$\begin{aligned} & \{J_{opt}, basis_{opt}\} \\ &= \arg \min_{J', basis} \sum_{k=1}^K \sum_{i=1}^n \left(\hat{y}_{J', \sigma_k, \hat{\lambda}, i}(x_{\sigma_k(i)}) - y(x_{\sigma_k(i)}) \right)^2, \end{aligned} \quad (39)$$

where $\hat{\lambda}$ is an estimator for the tuning parameter and $basis$ is a variable varying over the possible orthonormal bases for subspaces of $\langle \{e_1, \dots, e_J\} \rangle$.

The idea is that (39) is a sort of variance-bias expression that can be minimized to find the right number of basis elements. Minimizing in (39) means we are preventing the number of basis elements from being too small (high bias) or too large (high variance). This is embedded in (39) because it uses K independent orderings of the data and sequential predictive error (since we want $\hat{y}_{J', \sigma_k, \hat{\lambda}, i}(x_{\sigma_k(i)})$ at stage $\sigma_k(i)$ to use only the data $x_{\sigma_k(1)}, \dots, x_{\sigma_k(i-1)}$). Averaging over the permutations of the data points as they appear in the sequential predictive error means that it is reasonable to regard the empirical optimum as close to an actual optimum if it exists. In fact, in \mathcal{M} -open settings, it does not make sense to take limits of (39), so it is hard to prove theory. What is feasible is to seek a $\{J_{opt}, basis_{opt}\}$ that makes (39) small. This can be done numerically by a stochastic search provided we have a way to propose basis vectors.

A simplification of (39) is to replace the sequential prediction with some form of cross-validation; this saves on computing time.

Our method for data-driven random generation of basis elements is simple. Draw J bootstrap samples from the data of size n . For each bootstrap sample, define a basis element from the Nadaraya-Watson estimator. This gives J candidates for basis vectors. Next, apply Gram-Schmidt orthonormalization to form an orthonormal basis. In some cases, two or more of these basis elements may be so close as to be de facto the same. When this occurs, we reject the results and repeat the procedure from the beginning until we get J orthonormal basis elements. Note that any technique for nonparametric regression can be used in place of Nadaraya-Watson. In Sec. 5 we also use Gaussian process priors, see Rasmussen and Williams (2006), to generate function estimates that can be used as basis elements.

In addition, we may save on computing time by avoiding having to permute over basis elements if they have a natural ordering e.g., Fourier bases are ordered by frequency, Legendre polynomials are ordered by degree etc. Although there are many ways to order randomly generated basis elements, here we order them by using the size of regions in real spaces of dimension $\dim(x)$ as follows. Given the randomly generated orthonormal basis $\{\hat{e}_1, \dots, \hat{e}_J\}$ we first form the

nonparametric regression function $\hat{f}_{\hat{\lambda}} = \hat{f}_{\hat{\lambda}}(\cdot \mid (x_1, y_1), \dots, (x_n, y_n))$ using all the data. Then, we define the ordered basis $[e_1^*, \dots, e_J^*]$ by the criterion

$$|SA(\hat{f}_{\hat{\lambda}}) - SA(e_1^*)| \leq \dots \leq |SA(\hat{f}_{\hat{\lambda}}) - SA(e_J^*)|,$$

where $SA(f)$ is the surface area of the function f on domain, assuming that the domain is the same for all nonparametric regression estimators and compact. In the special case that x is unidimensional, $SA(\hat{f}_{\hat{\lambda}})$ is just the arc length over its compact domain in \mathbb{R} and when x is two-dimensional, $SA(\hat{f}_{\hat{\lambda}})$ is the area of the surface defined by $\hat{f}_{\hat{\lambda}}$ over its compact domain in \mathbb{R}^2 . Essentially, we are ordering the random basis elements by how close they are to a full data function estimator in terms of volume in $\mathbb{R}^{\dim x}$.

5. Computed examples

In this section we apply our technique described in previous sections to one \mathcal{M} -complete data set and one \mathcal{M} -open data set. The first is a ‘canned’ data set that is recognized to be difficult. The second is a new data set on soil moisture graciously provided by Prof. T. Franz, see Franz et al. (2014).

5.1. Forest Fires data

Consider the Forest Fires data set publicly available from the UC Irvine Machine Learning Repository. The sample size is $n = 517$ and there are 8 non-trivial explanatory variables related to the severity of a forest fire. The dependent variable is the burn area of the fire. Details and references can be found at <http://archive.ics.uci.edu/ml/datasets/Forest+Fires>. We regard the Forest Fires data set as \mathcal{M} -complete because a forest fire is a chemical reaction with a lot of randomness that cannot be quantified well. That is, there is so little about the process that is stable that it is unclear there is anything to estimate. However, it is plausible that there is a model, necessarily highly complex, that might accurately encapsulate the behavior of forest fires under a variety of environmental conditions. Of course, such a model could be so

complex that even though the data generator is \mathcal{M} -complete it is nearly \mathcal{M} -open. Thus, generating predictions may be the most appropriate approach even they have a large variability.

For our analysis, we divide the data randomly into two subsets, one for training and one for validation. The training set contains $n_1 = 267$ data points and the validation set contains $n_2 = 250$ data points. To generate a predictor, we assume the predictive analog of an additive model. That is, we form eight univariate models, each using one of the explanatory variables, and then we stack the predictors they generate. To generate each univariate model, we generated data-driven basis elements to use in a linear model. Each linear model has an associated point predictor and these are weighted by their stacking coefficients.

We consider two classes of data-driven basis elements. The first class is generated as discussed in Subsec. 4.2 using the Nadaraya-Watson estimator found using the `npreg()` function in R. The second class is generated using Gaussian process priors found using the `gausspr()` function in R. In both cases we used the default settings for the R functions and we set the number of basis elements to find to be $J = 10$. Unsurprisingly, J_{opt} assumed values and the high end of its range, 8, 9, and 10, but most often 10. Note that the value of J_{opt} depends on which variable was being used and this process is independent of the m in the ‘sum to m ’ constraint.

We summarize our results in Fig. 1. The left hand panel in panel shows that when the Nadaraya-Watson estimator is used to generate basis elements, then the optimal value of the constraint m is $m_{opt} = .91$. By contrast, the right hand panel shows that when the Gaussian process prior is used, we find $m_{opt} = 1.28$. In this case generating basis elements by Nadaraya-Watson gives a better prediction, i.e. a much lower cumulative predictive error than the Gaussian process prior method. Since we used the same J for the two estimators, we interpret our results in terms of a counterfactual: the Nadaraya-Watson method leads to basis elements that generate subspaces that are closer to the subspace containing what the true function would be if it existed. Note also that both optimal values of m are meaningfully different from one at least in




Figure 1: Plots of cumulative predictive error of stacking eight univariate predictors vs. m , the value of the constraint for the Forest Fires data. Left: Basis elements generated by Nadaraya-Watson. Right: Basis elements generated using Gaussian process priors with a radial basis function kernel.

terms of the predictive error they give.

5.2. Soil moisture data

As an example that is \mathcal{M} -open, we consider the Soil Moisture data set. The response variable is an interpolated form of the moisture in the topsoil. There are six explanatory variables three for location (two for location on a grid, one for elevation), two for soil electrical resistivity, and one for a standard ‘wetness index’ that is a function of elevation; see Franz et al. (2014) for a detailed description. The actual sample size is 18973 but for computational convenience, we randomly selected $n = 1000$ data points, dividing them into two sets of size 500, at random, for training and one for validation as before. We continued to set $J = 10$ and used Nadaraya-Watson and Gaussian process priors to generate basis elements. This time we used a polynomial kernel in the Gaussian process prior because other kernels did not permit convergence or gave the same result. Again, we found $J_{opt} = 8, 9, 10$.

The predictive error from stacking six univariate predictors for a range of




Figure 2: Plots of cumulative predictive error of stacking eight univariate predictors vs. m , the value of the constraint for the **Soil Moisture** data. Left: Basis elements generated by Nadaraya-Watson. Right: Basis elements generated using Gaussian process priors with a polynomial kernel.

m for basis elements generated using Nadaraya-Watson and Gaussian process priors are shown in Fig. 2. It is seen that the predictive performance is nearly the same for both cases and, in particular, the optimal predictive errors are small and the optimal constraint is nearly the same, $m_{opt} \approx 1$. Comparing with the results for the **Forest Fires** data, we see that the better the predictive performance is, the closer to one m_{opt} is and that when the predictive performance is weaker values of m_{opt} can be either larger or smaller than one. We suggest that had we chosen a smaller subset of the **Soil Moisture** data we would have found worse predictive performance and an m_{opt} further from one.

6. Discussion

Here we have formally established that leave-one-out cross-validation is asymptotically the optimal action in posterior risk for a variety of loss functions. We have used this to justify the coefficients in a stacking predictor since they are based on a cross-validation criterion. Stacking is a model averaging technique for prediction most effective when a true model is unavailable or may not even

exist. We have investigated theoretically and computationally the effect of different choices of constraints on the coefficients of the stacking predictor and suggest that not imposing any leads to the best result in the sense of minimizing predictive error. In fact, our examples suggest that a ‘sum to one’ constraint naturally emerges when the predictive error is small. We comment that obvious extensions of our technique of proof show that leave- k -out cross-validation by also be regarded as Bayes actions.

When the concept of a true model is problematic, it is natural to fall back on predictive methods. Indeed, it is possible that seeking a good predictor may be more useful than modeling when the model is very complex. For instance, if no simplification of the true model can be readily identified a model average predictor may give better performance in a mean squared error sense. This seems to be the case for our two examples here.

Finally, we recall that Stone (1977) showed that the Akaike information criterion (AIC) is asymptotically equivalent to leave-one-out cross-validation and that Shao (1997) shows these further asymptotically equivalent to the Mallows’ C_p criterion, the generalized cross-validation, and the ‘ GIC_2 ’ criterion. The implication from our main theorem here is that all of these methods can also be regarded as asymptotically Bayes optimal.

Appendix

Proof of the example in Section 1

Consider the two models $M_1 : Y = x_1\beta_1 + \epsilon$ and $M_2 : Y = x_2\beta_2 + \epsilon$ where the explanatory variables are orthogonal i.e., $x_1'x_2 = 0$, $E(\epsilon) = 0$, and $Var(\epsilon) = \sigma^2$. If we stack these two models with the sum to one constraint, by Corollary 3.1, we have

$$\begin{aligned}\hat{w}_1 &\propto \sum_{i=1}^n \hat{e}_{i,2}^2 - \hat{e}_1' \hat{e}_2 = \sum_{i=1}^n \frac{e_{i,2}^2}{(1 - h_{ii,2})^2} - \hat{e}_1' \hat{e}_2, \\ \hat{w}_2 &\propto \sum_{i=1}^n \hat{e}_{i,1}^2 - \hat{e}_1' \hat{e}_2 = \sum_{i=1}^n \frac{e_{i,1}^2}{(1 - h_{ii,1})^2} - \hat{e}_1' \hat{e}_2,\end{aligned}\tag{.1}$$

where $e_{i,j}$ and $h_{ii,j}$ are the ordinary residual and the leverage for case i under model M_j , respectively.

Since $h_{ii,j} = x_{ij}^2 / \sum_{i=1}^n x_{ij}^2$, if $h_{ii,j} \rightarrow 0$ as $n \rightarrow \infty$, then (.1) becomes

$$\begin{aligned}\hat{w}_1 &\propto \sum_{i=1}^n e_{i,2}^2 - \hat{e}'_1 \hat{e}_2 = (n-1)\sigma^2 - \hat{e}'_1 \hat{e}_2, \\ \hat{w}_2 &\propto \sum_{i=1}^n e_{i,1}^2 - \hat{e}'_1 \hat{e}_2 = (n-1)\sigma^2 - \hat{e}'_1 \hat{e}_2.\end{aligned}$$

Combining with the sum to one constraint, this yields $\hat{w}_1 = \hat{w}_2 = 1/2$.

Now, if we stack M_1 and M_2 with a sum to two constraint, then from Theorem 3.1 and similar arguments as above we have

$$\begin{aligned}\hat{w}_1 &\propto (n-1)\sigma^2 + \sum_{i=1}^n \hat{y}_{2,(-i)}^2 - \hat{e}'_1 \hat{e}_2 - \sum_{i=1}^n \hat{y}_{1,(-i)} \hat{y}_{2,(-i)}, \\ \hat{w}_2 &\propto (n-1)\sigma^2 + \sum_{i=1}^n \hat{y}_{1,(-i)}^2 - \hat{e}'_1 \hat{e}_2 - \sum_{i=1}^n \hat{y}_{1,(-i)} \hat{y}_{2,(-i)}.\end{aligned}$$

So, if $\sum_{i=1}^n \hat{y}_{1,(-i)}^2 = \sum_{i=1}^n \hat{y}_{2,(-i)}^2$ then combining with the sum to two constraint we get the weights now $\hat{w}_1 = \hat{w}_2 = 1$.

Proof of Theorem 4.1

Without loss of generality, assume \mathcal{M} and \mathcal{M}' are bases of $\langle \mathcal{M} \rangle = \langle \mathcal{M}' \rangle$ and hence $J = J'$.

For the basis \mathcal{M} , we have the decomposition

$$y = y_1 + y_2, \tag{.2}$$

where $y_1 = \sum_{j=1}^J \alpha_j f_j$, $y_2 = \sum_{j>J} \langle y, e_j \rangle e_j$, and $\{e_1, e_2, \dots\}$ is an orthonormal

basis for \mathcal{H} . Then, Theorem 3.3 gives

$$\begin{aligned}
Q_{\min}^{\mathcal{M}} &= \left\| (y(x_1), \dots, y(x_n))' - \sum_{j=1}^J \hat{w}_j(f_j(x_1), \dots, f_j(x_n))' \right\|^2 \\
&= \left\| \left(\sum_{j=1}^J \alpha_j f_j(x_1) + \sum_{j>J} \langle y, e_j \rangle e_j(x_1), \right. \right. \\
&\quad \left. \dots, \sum_{j=1}^J \alpha_j f_j(x_n) + \sum_{j>J} \langle y, e_j \rangle e_j(x_n) \right)' \\
&\quad \left. - \sum_{j=1}^J \hat{w}_j(f_j(x_1), \dots, f_j(x_n))' \right\|^2 \\
&= \left\| \sum_{j=1}^J (\alpha_j - \hat{w}_j)(f_j(x_1), \dots, f_j(x_n))' + \sum_{j>J} \langle y, e_j \rangle (e_j(x_1), \dots, e_j(x_n))' \right\|^2.
\end{aligned}$$

Since $f_j \in \langle \{e_1, \dots, e_J\} \rangle$ for $j = 1, \dots, J$ and $e_j \in \langle \{e_1, \dots, e_J\} \rangle^\perp$ for $j > J$, then

$$\begin{aligned}
Q_{\min}^{\mathcal{M}} &= \left\| \sum_{j=1}^J (\alpha_j - \hat{w}_j)(f_j(x_1), \dots, f_j(x_n))' \right\|^2 \\
&\quad + \left\| \sum_{j>J} \langle y, e_j \rangle (e_j(x_1), \dots, e_j(x_n))' \right\|^2 \\
&= \left\| \sum_{j=1}^J (\alpha_j - \hat{w}_j)(f_j(x_1), \dots, f_j(x_n))' \right\|^2 + \|y_2\|^2.
\end{aligned} \tag{.3}$$

Now, from (.2),

$$\begin{aligned}
\langle y, f_l \rangle &= \sum_{j=1}^J \alpha_j \langle f_j, f_l \rangle + \sum_{j=J+1}^n \langle y, e_j \rangle \langle e_j, f_l \rangle \\
&= \sum_{j=1}^J \alpha_j \langle f_j, f_l \rangle,
\end{aligned}$$

for $l = 1, \dots, J$. Therefore,

$$\alpha = T^{-1}c = \hat{w},$$

where T and c are given in Theorem 3.3, and hence (.3) yields

$$Q_{\min}^{\mathcal{M}} = \|y_2\|^2.$$

This result does not depend on \mathcal{M} ; therefore, $Q_{\min}^{\mathcal{M}} = Q_{\min}^{\mathcal{M}'}$.

Proof of Theorem 4.2

Without loss of generality, assume \mathcal{M} is orthonormal. Then, for \mathcal{M} , as in the proof of Theorem 4.1, we have

$$Q_{\min}^{\mathcal{M}} = \|y_2\|^2 = \sum_{j \geq J} \langle y, e_j \rangle^2. \quad (.4)$$

For \mathcal{N} , we now have the decomposition

$$y = \sum_{j=1}^{J-1} \langle y, f_j \rangle f_j + \sum_{j \geq J} \langle y, e_j \rangle e_j.$$

Then, from Theorem 3.3,

$$\begin{aligned} Q_{\min}^{\mathcal{N}} &= \left\| y - \sum_{j=1}^{J-1} \hat{w}_j f_j \right\|^2 \\ &= \left\| \sum_{j=1}^{J-1} \langle y, f_j \rangle f_j + \sum_{j \geq J} \langle y, e_j \rangle e_j - \sum_{j=1}^{J-1} \hat{w}_j f_j \right\|^2 \\ &= \left\| \sum_{j=1}^{J-1} (\langle y, f_j \rangle - \hat{w}_j) f_j + \sum_{j \geq J} \langle y, e_j \rangle e_j \right\|^2 \\ &= \left\| \sum_{j=1}^{J-1} (\langle y, f_j \rangle - \hat{w}_j) f_j \right\|^2 + \sum_{j \geq J} \langle y, e_j \rangle^2. \end{aligned} \quad (.5)$$

The desired inequality in Theorem 4.2 is obtained from (.4) and (.5).

References

- Bernardo, J., Smith, A., 2000. Bayesian Theory. John Wiley & Sons, Chichester.
- Billingsley, P., 2012. Probability and Measure. Wiley, New Jersey.

- Breiman, L., 1996. Stacked regressions. *Machine Learning* 24, 49–64.
- Clarke, B., 2003. Bayes model averaging and stacking when model approximation error cannot be ignored. *Journal of Machine Learning Research* , 683–712.
- Clyde, M., 2012. Bayesian perspectives on combining models. Slides from presentation at ISBA Kyoto , 648–649.
- Clyde, M., Iversen, E., 2013. Bayesian model averaging in the M-open framework, in: Damien, P., Dellaportas, P., Polson, N., Stephens, D. (Eds.), *Bayesian Theory and Applications*, Oxford University Press, Oxford. pp. 484–498.
- Franz, T., Wang, T., Avery, W., Finkenbiner, C., Brocca, L., 2014. Spatio-temporal characterization of soil moisture fields using cosmic ray neutron probes and data fusion. submitted .
- van de Geer, S., 2014. On the uniform convergence of empirical norms and inner products, with application to causal inference. *Elec. J. of Statistics* 8, 543–574.
- Nadaraya, E.A., 1964. On estimating regression. *Theory of Probability and its Applications* 9, 141–142.
- Ozay, M., Vural, F.T.Y., 2012. A new fuzzy stacked generalization technique and analysis of its performance. *IEEE* .
- Rasmussen, C.E., Williams, C.K.I., 2006. *Gaussian Processes for Machine Learning*. The MIT Press, Massachusetts.
- Rokach, L., 2010. Ensemble-based classifiers. *Artificial Intelligence Review* 33, 1–39.
- Shao, J., 1997. An asymptotic theory for linear model selection. *Statistica Sinica* 7, 221–264.

- Sill, J., Takacs, G., Mackey, L., Lin, D., 2009. Feature-weighted linear stacking. arxiv.org/pdf/0911.0460 .
- Smyth, P., Wolpert, D., 1999. Linearly combining density estimators via stacking. *Machine Learning Journal* 36, 59–83.
- Stone, M., 1977. Asymptotics for and against cross-validation. *Biometrika* 64, 29–38.
- Ting, K.M., Witten, I., 1999. Issues in stacked generalization. *Journal of Artificial Intelligent Research* 10, 271–289.
- Walker, S.G., Gutierrez-Pena, E., 1999. Robustifying bayesian procedures. *Bayesian Statistics* 6, 685–710.
- Watson, G.S., 1964. Smooth regression analysis. *The Indian Journal of Statistics, Series A* 26, 359–372.
- Wolpert, D., 1992. Stacked generalization. *Neural Networks* 5, 241–259.
- Wolpert, D., Macready, W., 1999. An efficient method to estimate bagging generalization error. *Machine Learning Journal* 35, 41–55.